Castle Laboratory

**Department of Operations Research and Financial Engineering**
**Princeton University**

Final Project Report:

# Information Acquisition and Representation Methods for Real-Time Asset Management

Principal Investigator:

Warren B. Powell
Princeton University
Department of Operations Research and Financial Engineering
Princeton, NJ 08544
powell@princeton.edu

Prepared for:

Air Force Office of Scientific Research
Grant: FA9550-05-1-0121

June 20, 2008

# REPORT DOCUMENTATION PAGE

*Form Approved*
*OMB No. 0704-0188*



**1. REPORT DATE** *(DD-MM-YYYY)*: 20-06-2008

**2. REPORT TYPE**: Final

**3. DATES COVERED** *(From - To)*: 01 FEB 2005 - 31 JAN 2008

**4. TITLE AND SUBTITLE**
INFORMATION AQUISITION AND REPRESENTATION METHODS FOR REAL-TIME ASSET MANAGEMENT

**5a. CONTRACT NUMBER**: FA9550-05-1-0121

**5b. GRANT NUMBER**: AFOSR

**5c. PROGRAM ELEMENT NUMBER**

**5d. PROJECT NUMBER**

**5e. TASK NUMBER**

**5f. WORK UNIT NUMBER**

**6. AUTHOR(S)**
Dr. Warren B. Powell

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
Princeton University
Department of Operations Research and Financial Engineering
Princeton, NJ 08544

**8. PERFORMING ORGANIZATION REPORT NUMBER**

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**
AFOSR/NL
875 Randolph
Arlington VA 22203

**10. SPONSOR/MONITOR'S ACRONYM(S)**: AFOSR

**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**: AFRL-SR-AR-TR-08-0386

**12. DISTRIBUTION/AVAILABILITY STATEMENT**
Distribution A

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**
Our research over the last six years has focused on the modeling and solution of complex resource allocation problems, which resulted in the integration of math programming within the emerging field of approximate dynamic programming, summarized in Warren B. Powell, Approximate Dynamic Programming: Solving the curses of dimensionality, John Wiley and Sons, New York, 2007. These techniques allow us to address problems which involve high-dimensional decision vectors (which routinely arise in resource allocation problems) under various forms of uncertainty. These techniques have proven themselves in a number of major project with civilian transportation companies (with implemented software), and several demonstration projects on air force logistics problems.

**15. SUBJECT TERMS**

| 16. SECURITY CLASSIFICATION OF: a. REPORT | b. ABSTRACT | c. THIS PAGE | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| | | | | | **19b. TELEPHONE NUMBER** *(Include area code)* |

Reset

## INSTRUCTIONS FOR COMPLETING SF 298

**1. REPORT DATE.** Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

**2. REPORT TYPE.** State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

**3. DATES COVERED.** Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

**4. TITLE.** Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

**5a. CONTRACT NUMBER.** Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

**5b. GRANT NUMBER.** Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

**5c. PROGRAM ELEMENT NUMBER.** Enter all program element numbers as they appear in the report, e.g. 61101A.

**5d. PROJECT NUMBER.** Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

**5e. TASK NUMBER.** Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

**5f. WORK UNIT NUMBER.** Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

**6. AUTHOR(S).** Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES).** Self-explanatory.

**8. PERFORMING ORGANIZATION REPORT NUMBER.** Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES).** Enter the name and address of the organization(s) financially responsible for and monitoring the work.

**10. SPONSOR/MONITOR'S ACRONYM(S).** Enter, if available, e.g. BRL, ARDEC, NADC.

**11. SPONSOR/MONITOR'S REPORT NUMBER(S).** Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

**12. DISTRIBUTION/AVAILABILITY STATEMENT.** Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

**13. SUPPLEMENTARY NOTES.** Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

**14. ABSTRACT.** A brief (approximately 200 words) factual summary of the most significant information.

**15. SUBJECT TERMS.** Key words or phrases identifying major concepts in the report.

**16. SECURITY CLASSIFICATION.** Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

**17. LIMITATION OF ABSTRACT.** This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Table of Contents

# 1. Summary of activities

Our research over the last six years has focused on the modeling and solution of complex resource allocation problems, which resulted in the integration of math programming within the emerging field of approximate dynamic programming, summarized in

> Warren B. Powell, *Approximate Dynamic Programming: Solving the curses of dimensionality*, John Wiley and Sons, New York, 2007.

These techniques allow us to address problems which involve high-dimensional decision vectors (which routinely arise in resource allocation problems) under various forms of uncertainty. These techniques have proven themselves in a number of major project with civilian transportation companies (with implemented software), and several demonstration projects on air force logistics problems.

Over the last three years, our research has focused on three related dimensions:

- Design and testing of ADP models and algorithms for specific applications arising both within the air force, and for civilian applications in freight transportation (these projects are funded by the companies). Military applications include a model for military airlift (where we study the value of information), and mid-air refueling.
- Theoretical issues in ADP – This work has focused on convergence proofs, optimal stepsizes to accelerate the rate of convergence, and theoretical work arising in the approximation of functions.
- Analysis of stochastic problems, including robustness, estimation of derivatives of stochastic optimization problems, stability of stochastic problems.
- Multiagent optimization and semicooperative control – This work addresses multiple decision makers and the analysis of different modes for working together.
- Optimal learning – This research arises in the ADP community under the name of “exploration vs. exploitation,” but has many applications outside of dynamic programming (e.g. sequential design of experiments, simulation optimization, and other stochastic optimization problems where probability distributions are unknown).

# 2. Approximate dynamic programming

Prior to our work, dynamic programming could be found in three major flavors:

- Markov decision processes – These models used discrete (flat) representations of states and actions, and were limited to extremely small problems.

- Reinforcement learning – This was the first field to develop the concept of solving dynamic programs under uncertainty, but this field has emphasized issues such as model-free dynamic programming (solving problems without an explicit mathematical model of the transition function), and approximating problems with large state spaces.
- Neuro-dynamic programming – This was the name given to approximation methods arising primarily in control theory with low-dimensional but continuous control vectors.

By contrast, the high-dimensional decision vectors that arise routinely in most resource allocation problems (such as those arising in air force logistics), were routinely modeled using the framework of deterministic math programming. These problems often resulted in large linear programs or, more often, large integer programs which were then solved using a variety of heuristics.

Approximate dynamic programming is often viewed as a set of techniques for solving dynamic programs that suffer from the well known "curse of dimensionality." In fact, there are three ways to view ADP, depending on the field that you are coming from:

- A decomposition method for very large-scale mathematical programs (including integer programs). This is how we first got into the field, and lately we have found that problems that were previously thought to be completely intractable (integer programs with millions of variables) can be solved using commercial solvers such as Cplex, without requiring the usual simplifying approximations that these models typically entail.
- An "optimizing-simulator" – It is extremely popular to tackle complex logistical problems (especially those involving uncertainty, but not necessarily) using simulation, where decisions are made over time. ADP is a method that can take a traditional simulator and turn it into a system where decision made at time t account for downstream impacts.
- Techniques for approximately solving dynamic programs – And of course, ADP does in fact provide a framework for producing high-quality solutions to dynamic programs that are otherwise computationally intractable.

Experienced experimentalists in math programming might view ADP as "just" a decomposition technique, but this strategy has allowed us to use a commercial solver such as Cplex to solve industrial-strength integer programs for companies as large as Yellow Freight System (the largest less-than-truckload carrier in the U.S.), Schneider National (one of the top three truckload carriers, managing over 15,000 drivers), Norfolk Southern Railroad (one of the top four North American railroads) and United Parcel Service. We are not able to produce globally optimal solutions, but the solutions are locally optimal (which means they look correct when examined up close), and do a near-optimal job in terms of making decisions over time (based on comparisons to optimal solutions when we can find them).

If someone has a simulator (popular in the military), ADP provides a framework for transitioning from myopic, rule-based decisions to decisions that balance here-and-now against the future within a rigorous mathematical framework.

Of course, ADP is a powerful algorithmic strategy for dynamic programs that are computationally intractable. Dynamic programming is widely dismissed because of the "curse of dimensionality." Our work identified three curses of dimensionality (the state space, the action space and the outcome space), and we describe how to use an important concept known as the post-decision state variable is critical for the merger of math programming (which handles high-dimensional decision vectors) and dynamic programming.

## 3. Theoretical research in ADP

ADP is a powerful modeling and algorithmic framework, but it is surprising how many open theoretical questions remain in the field. Theoretical research over the past three years has included the following:

- We have finally proved convergence for a forward approximate dynamic programming algorithm which involves a scalar controllable state variable with pure exploitation. This result required exploiting the concavity of the function in terms of the controllable state variable (common across all of our resource allocation problems). Prior theoretical work required assuming that we visit all states infinitely often. The state variable may include other dimensions that evolve exogenously.
- We also proved convergence for a "lagged asset acquisition problem" which arises when we make commitments against resources in the future.
- We developed a new stepsize algorithm that optimally adjusts to noise and bias. If the bias is zero, we obtain the optimal stepsize of 1/n. If the noise is zero, we obtain the optimal stepsize of 1.0.
- We derived a near-optimal method for combining statistics at multiple levels of aggregation.
- We have developed a new method for making expensive measurements known as the knowledge gradient algorithm. The algorithm is optimal for one measurement, and has been proven to be asymptotically optimal.
- We have recently developed a proof of convergence for an ADP algorithm using basis functions. Previously, the only general result was for an ADP algorithm using a fixed policy. Our result requires that the basis functions constitute a full basis.
- We have studied optimal learning policies in the context of the classic newsvendor problem, for which optimal learning results were limited to a single very special case. We show that the knowledge gradient policy closely matches the optimal policy on the special case, but can still be used for a wide range of problems.

# 4. Selected applications

We have demonstrated the ability of ADP to solve complex problems in both military and civilian settings. The civilian research (in transportation and logistics) all involves systems that were sponsored by and implemented at companies. Examples include:

A fleet simulator for Schneider National – This system required modeling the movements of approximately 6,000 drivers, each described by a 15-dimensional attribute vector. The system had to closely match historical performance at the company. We used ADP to optimize the flows, and found that this result produced the best results, outperforming a carefully engineered simulator. The system has been adopted by Schneider and is used extensively in a number of policy analyses.

Locomotive optimization at Norfolk Southern Railroad – The OR community has been trying for decades to develop an optimization model for locomotives. We now have the first production-quality locomotive optimization model. ADP allowed us to represent locomotives at a high level of detail. The model has been adopted by Norfolk Southern for planning purposes, and they intend to implement it as an operational tool.

Planning high-value spare parts for Embraer – The system has to plan inventories for over 400 high value spare parts, where the inventories are generally fairly small (often zero). In some cases, we might have only six spare parts spread among 18 locations. The inventories have to be planned to hit targets for total inventory cost and service, in the presence of random, spatially distributed demand.

# 5. Research reports sponsored by AFOSR

## *5.1. Journal articles (refereed)*

Frazier, P., W. B. Powell and S. Dayanik, “A Knowledge Gradient Policy for Sequential Information Collection,” Siam J. on Control and Optimization (to appear).

Simao, H. P., J. Day, A. George, T. Gifford, W. B. Powell, “An Approximate Dynamic Programming Algorithm for Large-Scale Fleet Management: A Case Application,” Transportation Science (to appear).

S. Dayanik, W. Powell, and K. Yamazaki (2007). “Index policies for discounted bandit problems with availability constraints,” Journal of Applied Probability, to appear.

A. Marar and W. B. Powell, “Capturing Incomplete Information in Resource Allocation Problems through Numerical Patterns,” European Journal of Operations Research, accepted June, 2008 (to appear).

Cheung, R. K.-M., N. Shi, W. B. Powell, and H. P. Simao, “An Attribute-Decision Model for Cross-Border Drayage Problem,” Transportation Research E: Logistics and Transportation Review, Volume 44, No. 2, pp. 217-234 (2008).

Topaloglu, H. and W.B. Powell, “Incorporating Pricing Decisions into the Stochastic Dynamic Fleet Management Problem,” Transportation Science, Vol. 41, No. 3, pp. 281-301 (2007).

Topaloglu, H. and W. B. Powell, “Sensitivity Analysis of a Dynamic Fleet Management Model Using Approximate Dynamic Programming” Operations Research, Vol. 55, No. 2, pp. 319-331 (2007).

Papadaki, K. and W. B. Powell, “Monotonicity in Multidimensional Markov Decision Processes for Batch Service Problems,” Operations Research Letters, Vol. 35, pp. 267-272 (2007).

George, A. and W. B. Powell, “Adaptive Stepsizes for Recursive Estimation with Applications in Approximate Dynamic Programming,” Machine Learning, Vol. 65, No. 1, pp. 167-198, (2006).

Dall’Orto, L. C., T. G. Crainic, J. E. Leal and W. B. Powell, “The Single-Node Dynamic Service Scheduling and Dispatching Problem,” *European Journal of Operations Research*, Vol. 170, No. 1, pp. 1-23 (2006).

Marar, A. W. B. Powell and S. Kulkarni, “Combining Cost-Based and Rule-Based Knowledge in Complex Resource Allocation Problems,” IIE transactions Vol. 38 (2), pp. 159-172 2006.

Topaloglu, H. and W.B. Powell, “Dynamic Programming Approximations for Stochastic, Time-Staged Integer Multicommodity Flow Problems,” *Informs Journal on Computing*, Vol. 18, No. 1, pp. 31-42 (2006).

Shapiro, J. and W.B. Powell, “A Metastrategy for Dynamic Resource Management Problems based on Informational Decomposition,” *Informs Journal on Computing*, Vol. 18, No. 1, pp. 43-60 (2006).

Topaloglu, H. and W.B. Powell, “A Distributed Decision-Making Structure for Dynamic Resource Allocation with Nonlinear Functional Approximations,” Operations Research, Vol. 53, No. 2, pp. 281-297 (2005)

## *5.2. Refereed book chapters and conference proceedings*

Kantor, P., P. Frazier and W. B. Powell, “Approximate Dynamic Programming in Knowledge Discovery for Rapid Response,” HICSS Conference, Hawaii, 2008.

Powell, W. B. and P. Frazier, “Approximate Dynamic Programming: Lessons from the field,” Invited tutorial, Proceedings, Winter Simulation Conference, 2008.

Powell, W. B. and P. Frazier, “Optimal Learning,” TutORials, INFORMS, 2008 (to appear).

Powell, W.B., “Real-time dispatching for truckload motor carriers,” in *Logistics Engineering Handbook* (G. Don Taylor, ed.), CRC Press, 2007, pp. 15-1 – 15-30.

Powell, W. B., “Approximate Dynamic Programming for High-Dimensional Problems,” Invited tutorial article for Winter Simulation Conference, December, 2007.

Powell, W.B., B. Bouzaiene-Ayari and H.P. Simao, “Dynamic Models for Freight Transportation,” Handbooks in Operations Research and Management Science: Transportation (G. Laporte and C. Barnhart, eds.), Elsevier, Amsterdam, 2007.

Nascimento, J. and W. B. Powell, “An Optimal ADP Algorithm for a High-Dimensional Stochastic Control Problem,” IEEE Conference on Approximate Dynamic Programming and Reinforcement Learning, April, 2007.

Frazier, P. and W. B. Powell, “The Knowledge Gradient Policy for Offline Learning with Independent Normal Rewards,” IEEE Conference on Approximate Dynamic Programming and Reinforcement Learning, April, 2007.

Powell, W.B. and H. Topaloglu, “Approximate Dynamic Programming for Large-Scale Resource Allocation,” in *Tutorials in Operations Research*, M. P. Johnson, B. Normal and Nicola Secomandi, eds. INFORMS, 2006.

Powell, W.B., “The Optimizing-Simulator: Merging Simulation and Optimization using Approximate Dynamic Programming,” Proceedings of the Winter Simulation Conference, December, 2005.

Powell, W.B., A. George, B. Bouzaiene-Ayari and H. Simao, “Approximate Dynamic Programming for High Dimensional Resource Allocation Problems,” Proceedings of the IJCNN, Montreal, August 2005.

Powell, W.B. and H. Topaloglu, “Fleet Management,” in *Applications of Stochastic Programming*, (S. Wallace and W. Ziemba, eds.), Math Programming Society - SIAM Series in Optimization, Philadelphia,pp. 185-216, 2005.

### *5.3. Books*

Powell, W.B., *Approximate Dynamic Programming for Asset Management*, John Wiley and Sons, 2007.

### *5.4. Doctoral dissertations*

The following doctoral dissertations were completed over the last three years.

Juliana Nascimento, 2008, “Approximate dynamic programming for complex storage problems,” McKinsey Consulting, Sao Paolo, Brazil

Gregory Godfrey,2007, “Nonlinear Approximation Method for Solving Stochastic, Dynamic Resource Allocation Problems,” First position: Metron Inc.

Abraham George, 2005, “Optimal Learning Strategies for Multi-Attribute Resource Allocation Problems,” First position: Research staff, Princeton University.

# 6. Personnel supported

Faculty:

- Professor Warren B. Powell

Professional staff:

- Dr. Hugo Simao
- Dr. Belgacem Bouzaiene-Ayari

Graduate students:

- Ilya Rhyzov (2nd year) – Ph.D.
- Peter Frazier (3rd year) – Ph.D.
- Juliana Nascimento – Ph.D.
- Abraham George - Ph.D.
- Dennis Panos (U.S. Navy) - MSE

# 7. Honors and awards

Finalist, Decision Analyais Society Student Paper Prize, "A Knowledge Gradient Policy for Sequential Information Collection," November, 2007.

Best paper prize at ICPR Americas conference, June, 2008, for: Simao, H. P. and W. B. Powell, "Approximate Dynamic Programming for Managing High Value Spare Parts."

# 8. Interactions/transitions

## *8.1. Participation/presentations at meetings, conferences, etc.*

### 8.1.1. Invited talks:

"Approximate Dynamic Programming: Solving the Curses of Dimensionality," Invited plenary speaker, ICPR Americas, Sao Paulo, Brazil, June 6, 2008.

"The Optimizing-Simulator for Capturing Real-World Military Operations," Air Mobility Command, Scott AFB, May 27, 2008.

"Approximate Dynamic Programming: Solving the Curses of Dimensionality," Invited plenary speaker, CIRRELT Workshop, Quebec City, May, 2008.

“Information collection and learning for nuclear detection,” Rutgers University, April, 2008.

“Approximate Dynamic Programming for High-Dimensional Problems,” Boston University, February 29, 2008.

Invited tutorial: “Approximate Dynamic Programming for Intelligent Simulation,” Winter Simulation Conference, Washington, D.C., 2007.

“Modeling control centers: Using approximate dynamic programming to model collective intelligence,” CHARRD Workshop, Department of Mechanical and Aeronautical Engineering, Princeton University, November, 2007.

“So You Want to get Funding From Industry”, Future Academicians Colloquium, Informs, Seattle, November, 2007.

Workshop on Modeling the National Ignition Facility, Lawrence Livermore National Laboratory, October 31, 2007.

“Dynamic Sensor Management,” Workshop on Nuclear Detection, Rutgers University, October 19, 2007.

“The Dynamic Energy Resource Model,” Energy Workshop, Lawrence Livermore National Laboratories, September, 2007.

Invited tutorial: “Approximate dynamic programming for high-dimensional applications,” Lawrence Livermore National Laboratory, Livermore, CA, July, 2007.

“Approximate Dynamic Programming for High-Dimensional Problems,” ExxonMobil, New Jersey, May 6 2007.

“Approximate Dynamic Programming for High-Dimensional Problems,” Department of Management Science and Engineering, Stanford University, May 4, 2007.

Invited tutorial: “Approximate Dynamic Programming: Solving problems the way people do,” Informs Practice Meeting, Vancouver, April 20, 2007.

Invited tutorial: “Approximate Dynamic Programming for High-Dimensional Problems,” IEEE Workshop on Approximate Dynamic Programming and Reinforcement Learning, Honolulu, Hawaii, April, 2007.

Distinguished UTC Seminar speaker: “Approximate Dynamic Programming for High-Dimensional Problems,” University of California at Davis, February, 2007.

“Approximate dynamic programming for High-Dimensional Resource Allocation,” University of Michigan, Department of Operations and Industrial Engineering, Ann Arbor, January, 2007.

“Approximate dynamic programming for military applications,” Joint mathematics meeting, New Orleans, January, 2007.

“From stochastic optimization to housing permits: Navigating the pot-holes to implementation in the real-world,” Seventh New Jersey Universities Homeland Security Research Consortium Symposium, Rutgers University, November 20, 2006.

Invited Tutorial:” Approximate Dynamic Programming for Large-Scale Resource Allocation,” Informs Annual Meeting, Pittsburgh, November, 2006.

“Tutorial: Approximate Dynamic Programming in Transportation and Logistics,” Workshop on Stochastics in Transportation and Logistics, Molde, Norway, June, 2006.

“Approximate Dynamic Programming for Solving High Dimensional Resource Allocation Problems in Transportation and Logistics,” DIMACS Workshop, ExxonMobil Research Center, New Jersey, April, 2006.

“Approximate Dynamic Programming for the Car Distribution Problem,” DIMACS Workshop, ExxonMobil Research Center, New Jersey, April, 2006.

“Tutorial: Modeling dynamic programs,” NSF Workshop and Outreach Tutorials on Approximate Dynamic Programming, April, 2006.

“Merging machine learning and math programming for solving high-dimensional resource allocation problems,” NSF Workshop and Outreach Tutorials on Approximate Dynamic Programming, April, 2006.

“Approximate Dynamic Programming for High-Dimensional Asset Allocation,” University of Iowa, January, 2006.

“Computational Methods for High Dimensional Dynamic Programs for Discrete Resource Allocation,” PICASSO Seminar series, Department of Computer Science, Princeton University, December, 2005.

“Approximate Dynamic Programming: Solving the Curses of Dimensionality,” IMA Workshop, University College London, September, 2005.

“A Robust Modeling and Algorithmic Strategy for Discrete, Dynamic Resource Allocation Problems,” AFOSR Grantees conference, St. Louis, August, 2005.

Powell, W.B., A. George, B. Bouzaiene-Ayari and H. Simao, “Approximate Dynamic Programming for High Dimensional Resource Allocation Problems,” International Joint Conference of the Neural Network Society, Montreal, August 2005.

“Markov Decision Processes: AI vs. OR”, AFOSR Workshop on Decision-Making in Adversarial Domains, Washington, D.C., May, 2005.

“Approximate Dynamic Programming for High-Dimensional Asset Allocation Problems,” University of Wisconsin, February, 2005.

“The Optimizing-Simulator for Freight Transportation,” Transportation Research Board, Washington, D.C., January, 2005.

### 8.1.2. Conference presentations:

"Approximate Dynamic Programming for the Single Machine Scheduling Problem," ICPR Americas, Sao Paulo, Brazil, June, 2008 (with Debora Ronconi).

"Approximate Dynamic Programming for the Management of High Value Spare Parts," ICPR Americas, Sao Paulo, Brazil, June, 2008 (with Hugo Simao)..

"An Index Policy for the Discounted Bandit Problem with Availability Constraints," Informs Annual Meeting, Seattle, November, 2007. (with Kazu Yamazaki)

"An Approximate Dynamic Programming Approach to the R&D Portfolio Problem," Informs Annual Meeting, Seattle, November, 2007. (with Lauren Hannah)

"A Knowledge Gradient Policy for Sequential Bayesian Ranking and Selection," Informs Annual Meeting, Seattle, November, 2007. (with Peter Frazier)

"Approximate Dynamic Programming for a Spare Parts Problem: The Challenge of Rare Events," Informs Annual Meeting, Seattle, November, 2007. (with Hugo Simao)

"A Dynamic Model for the Mitigation of Transmission Failure Risk," Informs Annual Meeting, Seattle, November, 2007. (with Johannes Enders)

"An Optimal Dynamic Hedging Strategy for Jet Fuel Costs," Informs Annual Meeting, Seattle, November, 2007. (with J. Nascimento).

"Value Function Approximations for Multistage Linear Programs," Informs Annual Meeting, Seattle, November, 2007.

"Pricing in Freight Transportation," Informs Annual Meeting, Pittsburgh, November, 2006 (with H. Topaloglu).

"An Optimal Approximate Dynamic Programming Algorithm to a Mutual Fund Problem," Informs Annual Meeting, Pittsburgh, November, 2006 (with J. Nascimento).

"Computational Experimentation with Two-Stage Stochastic Programs", Informs National Meeting, San Francisco, November, 2005. (with H. Topaloglu, J. Higle and L. Zhao).

"Merging Stochastic Programming and Approximate Dynamic Programming for High Dimensional Problems", Informs National Meeting, San Francisco, November, 2005. (with A. George).

"An Optimal Learning Algorithm for Purchasing Assets Over Time", Informs National Meeting, San Francisco, November, 2005. (with J. Nascimento)

"Stochastic Optimization for an Aging Electric Power Infrastructure", Informs National Meeting, San Francisco, November, 2005. (with J. Enders)

"Incorporating Pricing Decisions into the Stochastic Dynamic Fleet Management Problem", Informs National Meeting, San Francisco, November, 2005. (with H. Topaloglu)

"Using Distributed Computation to Accelerate Optimizing Simulators" IFORS Meeting, Hawaii, July, 2005. (with Jeff Day, Hugo Simao).

“Optimal Stepsizes for Approximate Dynamic Programming,” Informs Computing Conference, Washington, D.C., January, 2005.

### *8.2. Consultative and advisory functions*

I interact from time to time with the analysis group at AMC, as well as AFRL in Rome, NY. I have discussed the use of optimal learning techniques for video image processing for a project at AFRL (with Bruce Suter), and I recently gave a tutorial on approximate dynamic programming with military applications to the analysis group at AMC.

### *8.3. Transitions*

Our transitions have occurred along three lines:

- Communication of ideas to the analysis group at AMC and staff at AFRL.
- Development of a mid-air refueling model xxx
- Direct implementation of ideas through projects with the corporate partners of CASTLE Lab. This is the major path by which we test our ideas in the field. Industrial projects during 2005-2008 included work with Schneider National (one of the three largest truckload motor carriers), Netjets (largest fractional jet operator), Norfolk Southern Railroad (one of four class I railroads in the U.S.), and Embraer (major manufacturer of regional jets).
- Making software available for download from the CASTLE Lab web site (the PILOTVIEW diagnostic library is now available).
- Licensing of software through local consulting firms for use in systems for their clients. CASTLE Lab has a relationship with Princeton Consultants, Inc. (www.princeton.com) which implements optimization and simulation models in transportation and logistics.

Specific transitions to the industrial partners of CASTLE Lab over the last three years include:

1. Transition: Optimizing simulator for fleet planning at Schneider National. We have calibrated a system that models the flows of approximately 5,000 drivers of different types. Schneider is interested in knowing what types of drivers are most valuable to the fleet (similar to AMC asking which aircraft types are most valuable). It is almost impossible to answer this question using “what if” analyses. Our logic produces, from one run, estimates of the gradients with respect to each type of driver.

   Recipient: Schneider National, the nation’s largest truckload motor carrier.

2. Transition: A model for short-term operational forecasting of freight cars using the optimizing simulator. The optimizing simulator framework can

handle multiple sources of uncertainty (customer demands, order locations, travel times, equipment failures).

Recipient: Norfolk Southern Railroad

3 Transition: Operational, tactical and strategic planning of locomotives. This system uses the optimizing simulator concept, and in particular makes heavy use of techniques for modeling incomplete information through low dimensional patterns. The system was recently approved for production at Norfolk Southern Railroad, making it the first successful production optimization model developed for operational use in North America.

Recipient:
Norfolk Southern Railroad, which uses the system both for strategic planning of the fleet size, and short-term tactical forecasting of surpluses and deficits.

6. Transition: We developed a system for optimizing high-value spare parts. This problem involves designing inventory policies for parts where the inventories are often zero (only a few locations will have even a single spare). We have to design policies for hundreds of spare parts, so that the aggregate inventory cost is below a certain level, and where we achieve specific targets on aggregate service.

Recipient: Embraer